arXiv:math/0511126v1 [math.OA] 7 Nov 2005

An Explicit Duality for Finite Groups

Robert A. Cohen and Martin E. Walter

Abstract. Using a "3 by 3 matrix trick" we previously showed that multiplication in a C^*-algebra $\mathcal{A}$, an algebraic structure, is determined by the geometry of the C^*-algebra of the 3 by 3 matrices with entries from $\mathcal{A}$, $M_3(\mathcal{A})$. As an application of this algebra-geometry duality we now construct an order theoretic based duality theory for all groups which are either locally compact abelian or finite. This construction generalizes the van Kampen–Pontriagin duality for locally compact abelian groups.

Introduction

For any locally compact group G, the convex, partially ordered semigroup $P(G)$ of continuous positive definite functions on G, is a complete invariant of the group. In the discussion below, we show how $P(G)$ can be used to recover the algebraic structure of G when G is finite or abelian. In the process we outline a duality theory for locally compact groups

Throughout the paper, $C^*(G)$ and $L^1(G)$ will be defined as in [**D**], 13.9.1. For a C^*-algebra $\mathcal{A}$, we will use $\mathcal{A}^+$ to denote the positive part of $\mathcal{A}$, also as in [**D**]. The identity operator of $\mathcal{A}$ will be denoted by I and $B(\mathcal{H})$ stands for the bounded linear operators on the Hilbert space $\mathcal{H}$. For $a, b \in \mathcal{H}$, the convex hull of a and b is written as $\text{co}(a, b)$ and the orthogonal complement of the vector $\xi \in \mathcal{H}$ is written as $\xi^\perp$. The notation $M_n(\mathcal{A})$ denotes the n by n matrices with entries from $\mathcal{A}$; e.g., $\mathcal{A} = C^*(G)$ and $\mathcal{A} = \mathbb{C}$ are important examples used in this paper. The notation $\text{diag}(\lambda_1, ..., \lambda_n)$ will be used to denote a diagonal n by n matrix with λ_k in the k^{th} diagonal entry. The symbol $\mathbb{Z}_n$ will be used to denote the integers modulo n.

1. Binary Product and Order Structure Duality

For any locally compact group G, the product structure of G is determined by the order structure of $M_3(C^*(G))$. The correspondence can be seen as a special case of the following theorem:

2000 *Mathematics Subject Classification.* Primary 43A35, 43A65 Secondary 43A40, 20D99, 06F99.

Key words and phrases. group duality, positive definite function.

THEOREM 1.1. *Let $\mathcal{A}$ be a C^*-algebra with unit e. Suppose a, b are unitary elements of $\mathcal{A}$, i.e., $a^*a = aa^* = b^*b = bb^* = e$. Then if $x \in \mathcal{A}$ we have:*

$$\begin{pmatrix} e & a & x \\ a^* & e & b \\ x^* & b^* & e \end{pmatrix} \in M_3(\mathcal{A})^+$$

if and only if $x = ab$.

PROOF. See [**W5**]. □

If the hypothesis of Theorem 1.1 is true and $Q := \begin{pmatrix} e & a & x \\ a^* & e & b \\ x^* & b^* & e \end{pmatrix} \in M_3(\mathcal{A})^+$, then $\frac{1}{3}Q$ is a projection in $M_3(C^*(G))$. Unitaries and projections play such significant roles in operator algebra theory that an interesting observation must be made. We traditionally identify a group element a with a unitary operator in a C^*-algebra. In the case of finite groups, we may identify a with a unitary matrix with entries from the complex numbers (this fact is useful for many results in this paper). However by canonically identifying the unitary operator a with the block matrix $P = \frac{1}{2}\begin{pmatrix} e & a \\ a^* & e \end{pmatrix}$ we see that we can also identify the group element a with a projection operator, P. Theorem 1.1 shows further that we can construct the product ab via concatenation of $\begin{pmatrix} e & a \\ a^* & e \end{pmatrix}$ with $\begin{pmatrix} e & b \\ b^* & e \end{pmatrix}$ to produce the positive matrix $\begin{pmatrix} e & a & ab \\ a^* & e & b \\ (ab)^* & b^* & e \end{pmatrix}$, where the product ab corresponds to the two by two submatrix $\begin{pmatrix} e & ab \\ (ab)^* & e \end{pmatrix}$.

2. Constructing a Dual of $P_1(G)$

We often consider the semigroup $P_1(G)$ consisting of positive definite functions on G satisfying $p(e) \leq 1$. Indeed, $P_1(G)$ also suffices as a complete invariant of the group G [**W4**]. Our goal is to define the "dual of $P_1(G)$." To this end we first consider certain morphisms of $P_1(G)$. Let $P(\Gamma_n)$ be defined to be $\{M_n(\mathbb{C})^+, \circ\}$, the n by n complex, positive definite matrices with Schur–Hadamard product, i.e., if $(p_{jk}), (q_{jk})$ are in $M_n(\mathbb{C})^+$, then $(p_{jk}) \circ (q_{jk}) = (p_{jk}q_{jk})$, the componentwise product. Note that this collection of matrices happens to be identifiable with the set of complex positive definite functions on the groupoid Γ_n, of n by n matrix units. Hence we introduce the notation $P(\Gamma_n)$, cf., [**W3**]. Given $n = 1, 2, \ldots$ consider maps $\varphi_n : P_1(G) \to P(\Gamma_n)$ which satisfy:

(I) $\varphi_n(\lambda p + (1-\lambda)q)) = \lambda\varphi_n(p) + (1-\lambda)\varphi_n(q)$, $p, q \in P_1(G)$, $0 \leq \lambda \leq 1$, i.e., φ_n is affine;

(II) $\varphi_n(pq) = \varphi_n(p) \circ \varphi_n(q)$, $p, q \in P_1(G)$, i.e., φ_n is multiplicative;

(III) For $p, q \in P_1(G)$, if $p \geq q$, then $\varphi_n(p) \geq \varphi_n(q)$, i.e., φ_n is order-preserving.

What form must a matrix $\varphi_n(p)$ have? From the theory of the Fourier–Stieltjes algebra, $B(G)$, of G, cf., [**E**] there are elements x_{jk} in the spectrum of this commutative Banach algebra such that $\varphi_n(p)_{jk} = p(x_{jk})$ for all $p \in P_1(G)$. There are two possibilities for x_{jk}. The first is that x_{jk} is non-singular, i.e., there is a $g_{jk} \in G$ such that $p(x_{jk}) = p(g_{jk})$ for all $p \in P_1(G)$. The second possibility is that $p(x_{jk}) = 0$

for all $p \in P_1(G)$ with compact support. Such x_{jk} are called *singular*. We thus impose a fourth condition.

(IV) For each choice of j and k, $1 \leq j,k \leq n$, $p \in P_1(G) \mapsto \varphi_n(p)_{jk} \in \mathbb{C}$ is nonvanishing on at least one $p \in P_1(G)$ with compact support.

DEFINITION 2.1. A map $\varphi_n : P_1(G) \to P(\Gamma_n)$ satisfying I, II, III, IV is called a non-singular n-morphism of $P_1(G)$, i.e., φ_n is affine, multiplicative, order preserving, and nonsingular.

REMARK 2.2. For fixed $n = 1, 2, \ldots$, a non-singular n-morphism of $P_1(G)$ is of the form $p \in P_1(G) \mapsto (p(g_{jk}))_{1\leq j,k\leq n} \in P(\Gamma_n)$ for some $g_{jk} \in G$, $1 \leq j,k \leq n$.

Using the enveloping C^*-algebra and W^*-algebra of G, cf., [**W1**] we can consider G as a collection of unitaries in a C^*-algebra and apply Theorem 1.1 to refine the above remark, to get

REMARK 2.3. If $\varphi_n : P_1(G) \to P(\Gamma_n)$ is a non-singular n-morphism, then we hope in general that there exists $g_1, g_2, \ldots, g_n \in G$ such that

$$\varphi_n(p) = (p(g_j^{-1} g_k))_{1\leq j,k\leq n} \in P(\Gamma_n) \text{ for all } p \in P_1(G).$$

This hope is realized below in the case when G is abelian or finite.

We can now define the dual of $P_1(G)$.

DEFINITION 2.4. The dual of $P_1(G)$ consists of those functions $f : P_1(G) \mapsto \mathbb{C}$ such that there exists a non-singular n-morphism, φ_n, satisfying $f(p) = \varphi_n(p)_{jk}$ for some fixed j,k, $1 \leq j,k \leq n$, and all $p \in P_1(G)$. As we will see, the φ_n can be seen as "multiplication table" morphisms.

It is easy to see that the dual of $P_1(G)$ as a topological space is homeomorphic to G, but how does one discover its group structure? Pick any non-singular n-morphism φ_n and suppose f and h are in the dual of $P_1(G)$ and $f(p) = \varphi_n(p)_{j_0 k_0}$, $h(p) = \varphi_n(p)_{r_0 s_0}$ for all $p \in P_1(G)$. Thus we may have the following situation if $j_0 < k_0$, $r_0 < s_0$, $j_0 < r_0$, $k_0 < s_0$:

$$\varphi_n(p) = \begin{pmatrix} p(e) & & & & & & & & \\ & \ddots & & & & & & & \\ & & p(e) & \cdots & f(p) & \cdots & fh(p) & & \\ & & & \ddots & & & & & \\ & & & & p(e) & \cdots & h(p) & & \\ & & & & & \ddots & \vdots & & \\ & & & & & & \ddots & & \\ & & & & & & & p(e) & \\ & & & & k_0 & & s_0 & & \end{pmatrix} \begin{matrix} \\ \\ j_0 \\ \\ r_0 \\ \\ \\ \\ \end{matrix}$$

The product fh should appear as the element in the j_0, s_0^{th} entry for φ_n, i.e., $(fh)(p) = \varphi_n(p)_{j_0 s_0}$ for all $p \in P_1(G)$.

This product will be well defined modulo a detail discussed in the next section and (by checking diagrams) associative. The identity is any diagonal entry of any non-singular n-morphism, and f^{-1} is found to satisfy $f^{-1}(p) = \varphi_n(p)_{kj}$ whenever $f(p) = \varphi_n(p)_{jk}$ for particular, $1 \leq j,k \leq n$, and all $p \in P_1(G)$. In the rest of this paper we outline rigorous proofs of the above construction when G is abelian or finite.

3. Recovering Product Structure of G with $P_1(G)$

As noted previously, we have that $P_1(G)$ is a complete invariant of the group G. It follows that, given group elements a and b, $P_1(G)$ can sufficiently determine the product of a and b, but how explicitly can we construct this product? When trying to answer this question we were led to consider the elements of $P_1(G)$ in conjunction with Theorem 1.1. Translating an exercise in [**K**], we establish a motivating lemma:

LEMMA 3.1. *If $[a_{jk}] \in M_n(\mathcal{A})^+$ where $\mathcal{A}$ is a unital C^*-algebra and p is any state on $\mathcal{A}$, i.e., p is positive definite and $p(e) = 1$, then $[p(a_{jk})] \in M_n(\mathbb{C})^+$.*

PROOF. See [**K**], p.884 . □

The semigroup $P(G)$ identifies with the positive forms on $C^*(G)$, c.f., [**D**], 13.4. Without losing generality, we consider the sub-semigroup $K(G)$ of $P_1(G)$ consisting of positive definite functions on G satisfying $p(e) = 1$. We call these functions *states* on G and call $K(G)$ the *state space* of G. Note that any state on G can be uniquely extended to a state on $C^*(G)$. Thus, by Lemma 3.1,

$$\begin{pmatrix} e & a & x \\ a^* & e & b \\ x^* & b^* & e \end{pmatrix} \in M_3(C^*(G))^+ \Rightarrow \begin{pmatrix} 1 & p(a) & p(x) \\ \overline{p(a)} & 1 & p(b) \\ \overline{p(x)} & \overline{p(b)} & 1 \end{pmatrix} \in M_3(\mathbb{C})^+ \ \forall \ p \in K(G).$$

The converse of the implication above is not true in the non-abelian case. This can be seen by noting that, given a and b in non-abelian G with $ab \neq ba$, for all $p \in K(G)$ we have

$$\begin{pmatrix} 1 & p(a) & p(ab) \\ \overline{p(a)} & 1 & p(b) \\ \overline{p(ab)} & \overline{p(b)} & 1 \end{pmatrix} \in M_3(\mathbb{C})^+ \quad \text{and} \quad \begin{pmatrix} 1 & p(a) & p(ba) \\ \overline{p(a)} & 1 & p(b) \\ \overline{p(ba)} & \overline{p(b)} & 1 \end{pmatrix} \in M_3(\mathbb{C})^+.$$

Why? Because any positive definite function on G is positive definite on G^{op} (the "opposed" group where the product is reversed). The matrices above are positive for all $p \in K(G)$ by Theorem 1.1 and Lemma 3.1 applied to $C^*(G)$ and $C^*(G^{op})$, respectively. However,

$$\begin{pmatrix} e & a & ba \\ a^* & e & b \\ (ba)^* & b^* & e \end{pmatrix} \notin M_3(C^*(G))^+.$$

While slightly disappointing, the observation just described does not dash all hopes of constructing a "dual" version of Theorem 1.1. In fact, when G is abelian we immediately arrive at a pleasing and insightful result. In the following theorem we also define an important class $\{Q_p\}$ of complex matrices indexed by $p \in K(G)$.

THEOREM 3.2. *Let G be a locally compact abelian group with $a, b, x \in G$. Then*

$$Q_p := \begin{pmatrix} 1 & p(a) & p(x) \\ \overline{p(a)} & 1 & p(b) \\ \overline{p(x)} & \overline{p(b)} & 1 \end{pmatrix} \in M_3(\mathbb{C})^+ \ \forall \ p \in K(G)$$

if and only if $x = ab$.

PROOF. Checking that the determinant of Q_p is positive leads to (after a little algebraic manipulation) a fundamental inequality:

$$|p(x) - p(a)p(b)|^2 \leq (1 - |p(a)|^2)(1 - |p(b)|^2) \tag{3.1}$$

If p is any character on G, then $|p(a)| = 1$. Thus by (3.1), since characters are multiplicative, we have $p(x) = p(a)p(b) = p(ab)$. Also, since characters separate the points of G, we have $x = ab$.

The other direction is trivial using Theorem 1.1 and Lemma 3.1. □

Returning to the nonabelian case, it follows from (3.1) that if $|p(a)| = 1$ and $Q_p \in M_3(\mathbb{C})^+$ then $p(x) = p(a)p(b)$. Recalling that $Q_p \in M_3(\mathbb{C})^+$ when $x = ab$ or $x = ba$; we also have by (3.1) that if $|p(a)| = 1$ for some p, then $p(a)p(b) = p(ab) = p(ba)$. Finally, assuming $Q_p \in M_3(\mathbb{C})^+$ for some $x \in G$, we get that $p(x) = p(a)p(b) = p(ab) = p(ba)$ whenever $|p(a)| = 1$; or equivalently, whenever $|p(b)| = 1$. This suggests that x might be ab or ba.

Definition 3.3. For a unitary element a of a unital C^*-algebra $\mathcal{A}$, a state p on $\mathcal{A}$ such that $|p(a)| = 1$ is called an *eigenstate* of a.

These eigenstates are fundamental to our study. Note that the determinant of Q_p is zero at any eigenstate of a. Also, as just noted, x agrees with ab and ba at any eigenstate of a or b. In other words, if a, b, and x are unitary complex matrices written with respect to a basis of eigenvectors of a, then we immediately see that x agrees with ab and ba on the diagonal.

We will now state our main result for a "dual" version of Theorem 1.1 for finite groups. Note that the following theorem generalizes Theorem 3.2 for finite groups in the sense that x must be a product of a and b.

Theorem 3.4. *Suppose G is a finite group and $a, b, x \in G$. Then*

$$Q_p = \begin{pmatrix} 1 & p(a) & p(x) \\ \overline{p(a)} & 1 & p(b) \\ \overline{p(x)} & \overline{p(b)} & 1 \end{pmatrix} \in M_3(\mathbb{C})^+ \ \forall \ p \in K(G) \tag{3.2}$$

if and only if $x = ab$ or $x = ba$.

Proof. The proof will be outlined in the discussion to follow. The reverse implication follows immediately from Theorem 1.1, Lemma 3.1, and by noting that we can identify $K(G)$ with $K(G^{op})$. The proof in the forward direction is somewhat lengthy, and the argument is broken into several cases. The process can be generally described by the following:

Key Steps

(1) Show $x \in \langle a, b\rangle$ (the subgroup of G generated by a and b). In fact, we immediately get $x = a^s b a^t$ for some natural numbers s and t.
(2) Show $s + t \equiv 1 \pmod{r}$ where r divides $|a|$ (the order of a) and is determined by the construction of a unitary representation of G. Thus $x = a^s b a^{1-s+qr}$ for some integer q.
(3) Show $s \equiv 0$ or $s \equiv 1 \pmod{|a|}$.
(4) Show $q = 0$, and thus conclude that $x = ab$ or $x = ba$.

We argue the case where $r = 1$ and $2 \leq r \leq |a|$ separately. We treat separately as well the cases $b \in Hb^{-1}H$ and $b \notin Hb^{-1}H$ where $H = \langle a\rangle$, the cyclic subgroup of G generated by a.

In the next section we include a full proof of Key Step (1) that illustrates a general approach used in the other steps. We also include a proof for the case when $ab = ba^m$ for some m and $b \in Hb^{-1}H$ to demonstrate the usefulness of the inequalities derived in the next section. In virtually all components of the proofs of the Key Steps above, we repeatedly appeal to these inequalities.

4. Fundamental Inequalities

Assuming that $\{Q_p : p \in K(G)\} \subset M_3(\mathbb{C})^+$ leads immediately to several inequalities involving a, b, and x, which are fixed. We define the real valued function $f(p) := \mathrm{Det}(Q_p)$ from $K(G)$ to $\mathbb{R}$. First, by assumption, $f(p)$ must be nonnegative, thus

$$f(p) = 1 + 2\,\mathrm{Re}[p(a)p(b)\overline{p(x)}] - |p(x)|^2 - |p(b)|^2 - |p(a)|^2 \geq 0 \ \forall \ p \in K(G). \tag{4.1}$$

We have already seen the inequality (4.1) rearranged to the aesthetically pleasing (3.1). We also noted that $f(p) = 0$ when p is an eigenstate of a, and thus by the assumption that $Q_p \in M_3(\mathbb{C})^+$ for all $p \in K(G)$, such eigenstates are minima of f.

We can always construct positive definite functions on G by finding a unitary representation π of G in $M_n(\mathbb{C})$ for some n and applying vector states, i.e., states of the form $p(\cdot) = \langle \pi(\cdot)\xi, \xi\rangle$ for some unit vector $\xi \in \mathbb{C}^n$. In fact, all positive definite functions on G arise in this manner [**D**], §13. Also, in this setting, eigenstates of a are just vector states corresponding to eigenvectors of $\pi(a)$. For generality, we develop a family of fundamental inequalities that hold in a $B(\mathcal{H})$ setting.

THEOREM 4.1. *Let $a, b, x \in B(\mathcal{H})$ where the operator a is unitary with unit eigenvector $\xi \in \mathcal{H}$ and corresponding eigenvalue λ. Suppose $Q_p \in M_3(\mathbb{C})^+$ for all vector states p on $B(\mathcal{H})$. Then for any unit vector $\eta \in \xi^\perp$,*

$$|\langle (x - ab)\eta, \xi\rangle + \langle (x - ba)\xi, \eta\rangle|^2 \leq (1 - |\langle b\xi, \xi\rangle|^2)\|(\lambda I - a)\eta\|^2. \tag{4.2}$$

Proof: Consider the smooth path $\xi_t = \cos(t)\xi + \sin(t)\eta$ through the unit sphere of $\mathcal{H}$ parameterized by $t \in \mathbb{R}$. Denoting the derivative of ξ_t by ξ_t', note that $\xi_0 = \xi$ and $\xi_0' = \eta$. Also, ξ_t induces a smooth path through the vector state space of $B(\mathcal{H})$ where p_t is just the vector state determined by ξ_t. Thus the function $f(p_t) = \mathrm{Det}(Q_{p_t})$ is a smooth map from $\mathbb{R}$ to $\mathbb{R}$, and classical calculus can be applied. Defining $\hat{a}$ as a function on $\mathbb{R}$ by $\hat{a}(t) = p_t(a)$ (similarly for $\hat{x}$ and $\hat{b}$) we can view f as:

$$f(t) = 1 + 2\,\mathrm{Re}[\hat{a}(t)\hat{b}(t)\overline{\hat{x}(t)}] - |\hat{a}(t)|^2 - |\hat{b}(t)|^2 - |\hat{x}(t)|^2. \tag{4.3}$$

Since $f(0) = 0$ is an assumed minimum, we have that $f''(0) \geq 0$. Recall that $\hat{a}(0) = \lambda$ and that $\hat{x}(0) = \lambda\hat{b}(0)$ since $p(x) = p(a)p(b)$ at eigenstates. Therefore, twice differentiating f leads to, after some manipulation

$$|\hat{x}'(0) - \lambda\hat{b}'(0)|^2 \leq (|\hat{b}(0)|^2 - 1)\,\mathrm{Re}[\overline{\lambda}\hat{a}''(0)]. \tag{4.4}$$

Applying Leibniz' rule to evaluate (4.4), we arrive at (4.2). □

We recognize that the family of inequalities captured in (4.2) can be sharpened.

COROLLARY 4.2. *Let $a, b, x \in B(\mathcal{H})$ satisfy the hypotheses of Lemma 4.1. Then for any vector $\eta \in \mathcal{H}$,*

$$|\langle (x - ab)\eta, \xi\rangle| + |\langle (x - ba)\xi, \eta\rangle| \leq \sqrt{1 - |\langle b\xi, \xi\rangle|^2}\|(\lambda I - a)\eta\|. \tag{4.5}$$

Proof: We can choose any $\eta \in \mathcal{H}$ by writing $\eta = \tilde{\eta} + c\xi$ for $\tilde{\eta} \in \xi^\perp$ and $c \in \mathbb{C}$ and applying the unit vector $\frac{e^{i\theta}\eta}{\|\eta\|}$ to (4.2) for an appropriate choice of θ to arrive at (4.5). □

Of particular interest to us is the case when $B(\mathcal{H}) = M_n(\mathbb{C})$, since finite groups always have finite dimensional representations as complex matrices. In particular, the following corollary holds when the elements a, b, and x of G are represented as

unitary matrices. In this case we can always write the matrices with respect to a basis such that a is diagonal.

COROLLARY 4.3. *Suppose $a, b, x \in M_n(\mathbb{C})$, where a is diagonal and unitary. We write $a = diag(\lambda_1, ..., \lambda_n)$, $b = [b_{jk}]$, and $x = [x_{jk}]$. Furthermore, suppose $Q_p \in M_3(\mathbb{C})^+$ for each vector state p on $M_n(\mathbb{C})$. Then for each $j \in \{1, ..., n\}$ we have*

$$(4.6)\quad |\sum_{k=1}^{n} r_k(e^{i\phi_k}(x_{jk}-\lambda_j b_{jk})+e^{-i\phi_k}(x_{kj}-\lambda_j b_{kj}))|^2 \leq (1-|b_{jj}|^2)\sum_{k=1}^{n} r_k^2|\lambda_j-\lambda_k|^2$$

where (4.6) holds for any choice of $r_k, \phi_k \in \mathbb{R}$, for $k = 1, ..., n$.

Proof: We arrive immediately at (4.6) by applying Lemma 4.1 where $\xi = e_j$, the j^{th} standard elementary basis vector for $\mathbb{C}^n$, and η is defined componentwise by $\eta_k = r_k e^{i\phi_k}$. Note that η may be arbitrary by Corollary 4.2. □

The family of inequalities captured in (4.6) are appealed to repeatedly in the proof of Theorem 3.4. Although various other inequalities have been derived pertaining to our problem, c.f., [**C**] Ch. 4, these have thus far been the most useful for our purposes. In the case of $M_n(\mathbb{C})$, the corresponding diagonal entries of x, ab, and ba when a is diagonal led us to look for constraints on the off-diagonal elements of x. The above inequalities clearly contain comparisons of x with ab and ba on the left hand side of (4.2). The terms $x_{jk} - \lambda_j b_{jk}$ and $x_{kj} - \lambda_j b_{kj}$ in (4.6) are, when $j \neq k$, the off-diagonal entries from the matrices $x - ab$ and $x - ba$, respectively (when a is a diagonal unitary matrix).

5. Utilizing a Cayley Representation

Throughout this section we assume G is a finite group of order $|G|$, and that $a, b, x \in G$. Also, we use n_a to denote the order of a, i.e., the smallest natural number satisfying $a^{n_a} = e$, and $\langle a \rangle$ to denote the cyclic subgroup of G generated by a. For the majority of the proof of Theorem 3.4 we make use of a Cayley representation π of G.

DEFINITION 5.1. If G is a finite group enumerated as $g_1, ..., g_{|G|}$ we form the matrix $[g_j^{-1}g_k] \in M_{|G|}(C^*(G))$. This matrix is sometimes called a *multiplication table* for G. Then associated with this enumeration we have a *Cayley representation* π of G defined by

$$\pi(g) = [g_{jk}], \text{ where } g_{jk} = \begin{cases} 1, & \text{if } g = g_j^{-1}g_k \\ \\ 0, & \text{if } g \neq g_j^{-1}g_k. \end{cases}$$

Thus, for each $g \in G$, $\pi(g)$ is a $|G| \times |G|$ permutation matrix, hence unitary. It is easily verified that π is a representation of G. Also, π is clearly faithful, and any vector state on $M_{|G|}(\mathbb{C})$ gives us a state on G. Another important feature of Cayley representations for our purposes is highlighted in the next lemma.

LEMMA 5.2. *Suppose G is a finite group with Cayley representation π, and $g, h \in G$. If $\pi(g)$ agrees with $\pi(h)$ in any nonzero entry then $g = h$.*

PROOF. This follows directly from the definition of a Cayley representation. □

With the generality of a Cayley representation and the sharpness of Lemma 5.2 in mind, we consider a strategic way to enumerate the group elements. The general

approach will be described in the proof of the following lemma, which is a proof of Key Step (1) above.

LEMMA 5.3. *Suppose G is a finite group and $a, b, x \in G$ satisfy $Q_p \in M_3(\mathbb{C})^+$ for all $p \in K(G)$. Then $x = a^s b a^t$ for some $s, t \in \mathbb{Z}_{n_a}$.*

PROOF. If $b \in \langle a \rangle$ or $x \in \langle a \rangle$ then the lemma follows by noting that any positive definite function p on a subgroup H of G can be extended to a positive definite function on G by simply letting $p(g) = 0$ for $g \notin H$. If p is any character on $H = \langle a \rangle$ extended in this fashion, then from (3.1) it follows that $p(x) = p(a)p(b) \neq 0$, implying that both $b \in \langle a \rangle$ and $x \in \langle a \rangle$. Therefore $p(x) = p(ab)$ for each character on $\langle a \rangle$, and hence $x = ab$. Also, if $a = e$ the lemma follows trivially from (3.1).

Thus assume $a \neq e$, $b \notin \langle a \rangle$, and $x \notin \langle a \rangle$. Let $H = \langle a \rangle$ and choose $\{\beta_1, ..., \beta_d\} \subset G$ with $\beta_1 = e$ and $\beta_2 = b$, such that $\bigcup_{k=1}^{d} \beta_k H$ is a coset decomposition of G. We enumerate G in such a way that the top row of the multiplication table reads

$$e, a, a^2, ..., a^{-1}, b, ba, ba^2, ..., ba^{-1}, \beta_3, \beta_3 a, ..., \beta_3 a^{-1}, ..., \beta_d, \beta_d a, ..., \beta_d a^{-1}.$$

or more concisely, if H is enumerated as $h_1 = e, h_2 = a, h_3 = a^2$, etc. we may write the top row in shorthand as

$$H, bH, \beta_3 H, ..., \beta_d H.$$

We define our Cayley representation π by this enumeration of G, and consider $\pi(a)$, $\pi(b)$, and $\pi(x)$ as $|G|$ by $|G|$ permutation matrices comprised of blocks of size n_a by n_a. Since π is faithful we will, for brevity, identify a with $\pi(a)$, etc., without ambiguity. We use the notation A_{jk}, B_{jk}, and X_{jk} to denote the j, k^{th} block of a, b, and x respectively. In this case, a is completely determined as a block diagonal matrix where each diagonal block has the form

$$A_{jj} = \begin{pmatrix} 0 & 1 & 0 & 0 & ... & 0 \\ 0 & 0 & 1 & 0 & ... & 0 \\ 0 & 0 & 0 & 1 & ... & 0 \\ . & . & . & . & ... & 0 \\ . & . & . & . & ... & 1 \\ 1 & 0 & 0 & 0 & .... & 0 \end{pmatrix}.$$

With this nice form it is easy to construct eigenvectors (and thus eigenstates) for a. Also, by assumption, b and x are matrices with all zeros in the diagonal blocks. We are led to consider what form b and x have on off-diagonal blocks, most importantly those in the first row and column of blocks. For the argument in this proof, we need only the upper two block by two block submatrices of a, b, and x, but it demonstrates a general construction and approach. We use the notation $\#(B_{jk})$ and $\#(X_{jk})$ to denote the number of nonzero entries in $B_{j,k}$ and $X_{j,k}$, respectively.

For any $\phi \in \mathbb{R}$ define the eigenvector $\xi \in \mathbb{C}^{|G|}$ of a componentwise by

$$\xi_j = \begin{cases} \frac{e^{i\phi}}{\sqrt{2n_a}}, & j = 1, ..., n_a \\ \frac{1}{\sqrt{2n_a}}, & j = n_a + 1, ..., 2n_a \\ 0, & j = 2n_a + 1, ..., |G|. \end{cases}$$

Let p be the eigenstate of a of the form $p(\cdot) = \langle \cdot \xi, \xi \rangle$. Since $p(a) = 1$, by (3.1) we have $p(x) = p(a)p(b) = p(b)$. By our construction this translates to

$$\frac{1}{2n_a}[\#(X_{12})e^{-i\phi} + \#(X_{21})e^{i\phi}] = p(x) = p(b) = \frac{1}{2n_a}[\#(B_{12})e^{-i\phi} + \#(B_{21})e^{i\phi}]$$

for any choice of ϕ. Therefore, by the linear independence of $e^{i\phi}$ and $e^{-i\phi}$, $\#(B_{12}) = \#(X_{12})$ and $\#(B_{21}) = \#(X_{21})$. Since there is at least one nonzero entry in B_{12}, e.g., $b_{1,n_a+1} = 1$, there is at least one nonzero entry in X_{12}. By our construction of the Cayley representation, this nonzero entry corresponds to an element of the double coset HbH. Thus, by Lemma 5.2, $x = a^s b a^t$ for some $s, t \in \mathbb{Z}_{n_a}$. □

As we go on to Key Steps (2) through (4), the arguments proceed in similar fashion, beginning with a carefully chosen enumeration of G and construction of the corresponding Cayley representation. We choose the enumeration so that a has the same block diagonal form as in the preceding proof and that b has an easily manipulatable form. Next, eigenvectors of a are constructed and utilized along with Lemma 4.6 to further restrict the form of x as a word in a and b. We take full advantage of the fact that the eigenvalues of a are the n_a^{th} roots of unity, which work nicely in the context of Lemma 4.6.

As the steps for the different cases proceed the choice of enumeration and construction of eigenvectors becomes unexpectedly more delicate and involved, see [**C**], but follow in the same spirit as described in the proof of Lemma 5.3. We include a proof for the particular case when the relationship $ab = ba^m$ holds for some m and $b \in Hb^{-1}H$. An example of a finite group for which this case pertains would be any of the quaternion groups. The purpose of including this proof is to give the reader a small sample of the usefulness of Theorem 4.1 and its corollaries.

THEOREM 5.4. *Suppose G is a finite group and $a, b, x \in G$ satisfy $Q_p \in M_3(\mathbb{C})^+$ for all $p \in K(G)$. Suppose further that $ab = ba^m$ for some $m \in \mathbb{Z}_{n_a}$ and that $b \in Hb^{-1}H$. Then $x = ab$ or $x = ba$.*

Proof: If $a = e$, $b \in \langle a \rangle$, or $x \in \langle a \rangle$ then $x = ab = ba$ by the first paragraph of the proof of Lemma 5.3. Thus assume $a \neq e$, $b \notin \langle a \rangle$, and $x \notin \langle a \rangle$.

We know by Lemma 5.3 that $x = a^s b a^t$ for some s and t. Therefore, since $ab = ba^m$, we have that $x = ba^l$ for some $l \in \mathbb{Z}_{n_a}$. We begin by showing that $l \equiv 1 + q(m-1) \pmod{n_a}$ for some $q \in \mathbb{Z}_{n_a}$.

Let π be the same Cayley representation of G used in the proof of Lemma 5.3 and suppose $\lambda = e^{2\pi i/(m-1,n_a)}$, where $(m-1, n_a)$ denotes the greatest common divisor of $m-1$ and n_a. Again, for brevity, identify a with $\pi(a)$, etc. Let the eigenvector ξ of a be defined componentwise by

$$\xi_j = \begin{cases} \frac{\lambda^j}{\sqrt{2n_a}} e^{i\phi}, & j = 1, 2, ..., n_a \\ \frac{\lambda^j}{\sqrt{2n_a}}, & j = n_a + 1, n_a + 2 ..., 2n_a \\ 0, & \text{otherwise.} \end{cases} \tag{5.1}$$

where ϕ is arbitrary. Then ξ is a unit eigenvector of a with eigenvalue λ. It follows, since $ab = ba^m$, that $b\xi$ is an eigenvector of a with eigenvalue $\lambda^m = \lambda$. By our constructions it follows that, for some $z \in \mathbb{C}$, we have $\langle b\xi, \xi \rangle = \frac{1}{2n_a}(n_a e^{-i\phi} + z e^{i\phi})$. We choose ϕ so that $\langle b\xi, \xi \rangle \neq 0$.

Recall that x must agree with ab at eigenstates of a. Thus $\lambda \langle b\xi, \xi \rangle = \langle ab\xi, \xi \rangle = \langle x\xi, \xi \rangle = \langle ba^l \xi, \xi \rangle = \lambda^l \langle b\xi, \xi \rangle$. In other words, we have $\lambda^l = \lambda$ and hence $l - 1 \equiv 0 \pmod{(m-1, n_a)}$. Therefore, $l = 1 + s(m-1, n_a)$ for some $s \in \mathbb{Z}$, giving $l \equiv 1 + q(m-1) \pmod{n_a}$ for some $q \in \mathbb{Z}_{n_a}$, since $(m-1, n_a)$ is congruent, mod n_a, to a multiple of $m-1$.

For uniqueness, we may assume $q \in \{0, 1, ..., |m-1|-1\}$ (where $|m-1|$ denotes the order of $(m-1)$ in $\mathbb{Z}_{n_a}$). The next step is to show that q must be 0 or 1.

Note that $ab = ba^m$ implies $a^k b = ba^{km}$ for each $k \in \mathbb{N}$ and hence $k \equiv 0 \pmod{n_a}$ if and only if $km \equiv 0 \pmod{n_a}$. Thus m has order n_a in $\mathbb{Z}_{n_a}$, and hence there is some unique $m^{-1} \in \mathbb{Z}_{n_a}$ such that $mm^{-1} \equiv 1 \pmod{n_a}$. It follows that $x = ba^l = a^{m^{-1}l}b$. Let λ be an arbitrary eigenvalue of a and define the eigenvector ξ of a as in (5.1). Letting $\eta = b\xi$ and writing $x = a^{m^{-1}l}b$ in the first term and $x = ba^l$ in the second term of (4.5), we have

$$|\lambda^{m^{-1}l} - \lambda||\langle b^2\xi, \xi\rangle| + |\lambda^l - \lambda| \leq |\lambda - \lambda^m|.$$

However, since $b \in Hb^{-1}H$ it follows that $\langle b^2\xi, \xi\rangle = \lambda^t$ for some $t \in \mathbb{Z}_{n_a}$. Also, since $l = 1 + q(m-1)$, the previous inequality may be written

$$|\lambda^{(1-q)(m^{-1}-1)} - 1| + |\lambda^{q(m-1)} - 1| \leq |\lambda^{m-1} - 1|.$$

Using the fact that $\{\lambda^{m-1} | \lambda \text{ is an eigenvalue of } a\}$ is a subgroup of the n_a^{th} roots of unity, we select λ so that λ^{m-1} is as close as possible (but not equal) to 1. Specifically we choose λ such that $\lambda^{m-1} = e^{2\pi i/|m-1|}$. Noting that $|\lambda^{q(m-1)} - 1| \leq |\lambda^{m-1} - 1|$ we can see that the only choices for q are 0, 1, and $|m-1| - 1$. If $q = |m-1| - 1$ then we must have $\lambda^{2(m^{-1}-1)} = 1$ for all eigenvalues λ of a. Since $|m-1| = |m^{-1} - 1|$, this implies that $2 = |m^{-1} - 1| = |m - 1|$, and thus the only choices for q are 0 or 1. Therefore, $q = 0$ or $q = 1$, implying $x = ba$ or $x = ba^m = ab$. □

Again, the previous proof only gives a glimpse of the technique used for each of the various cases. The reader may note that the proof needs only minor modifications to obtain the case where $b \notin Hb^{-1}H$. Refined versions of the proofs in [**C**] will appear shortly.

References

[C] R.A. Cohen, *Construction of an Order Theoretic Duality for Certain Groups*, Ph.D. Thesis, University of Colorado at Boulder, 2004.

[D] J. Dixmier, *Les C^*-algèbres et leurs représentations*, Cahiers Scientifiques, Fasc.29, Gauthier-Villars, Paris, 1964.

[E] P. Eymard, *L'algèbre de Fourier d'un groupe localement compact*, Bull. Soc. Math. France **92** (1964), 181–236.

[K] R.V. Kadison and J.R. Ringrose, *Fundamentals of the Theory of Operator Algebras* **2**, Academic Press, 1983.

[W1] M.E. Walter, *W^*-Algebras and nonabelian harmonic anaylysis*, J. Functional Analysis **11** (1972), 17–38.

[W2] M.E. Walter, *Duality Theory for Nonabelian Locally Compact Groups*, **XXII**, (1977), 47–59.

[W3] M.E. Walter, *C^*-Algebras and the Classification of Finite Groups*, Proc.of Symposia in Pure Mathematics, Volume **51**, Part 2, (1990), 345–354.

[W4] M.E. Walter, *Semigroups of Positive Definite Functions and Related Topics*, Harmonic Analysis and Hypergroups, Birkhäuser, 1998, 215–226.

[W5] M.E. Walter, *Algebraic Structures Determined by 3×3 Matrix Geometry*, Proceedings of the American Mathematical Society, **131** (2003), 2129–2131.

Department of Mathematics, University of Colorado, Campus Box 395, Boulder, Colorado, 80309-0395

E-mail address: robert.cohen@euclid.colorado.edu

Department of Mathematics, University of Colorado, Campus Box 395, Boulder, Colorado, 80309-0395

E-mail address: martin.walter@euclid.colorado.edu